

UCRL- 95911
PREPRINT

UNIFIED SUN AND SKY ILLUMINATION
FOR SHADOWS UNDER TREES

Nelson Max

CIRCULATING COPY
SUBJECT TO RECALL
IN TWO WEEKS

This paper was prepared for submittal to
ACM SIGGRAPH '87
Anaheim, CA, July 27-31, 1987

January 12, 1987

Lawrence
Livermore
National
Laboratory

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Unified Sun and Sky Illumination for Shadows under Trees

NELSON MAX

Lawrence Livermore National Laboratory

Abstract

Assume the sun, sky and clouds are represented as a raster image on a plane at infinity, and the shadows of a tree canopy are represented as a raster transparency mask on a canopy plane. Then the illumination on the ground is the correlation of these two rasters, which can be rapidly computed by Fast Fourier Transform techniques.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism — Color, shading, shadowing, and texture

General Terms: Computer graphics, image synthesis, shading, shadows

Additional Key Words and Phrases: penumbra, sky illumination, correlation, Fast Fourier Transform

Introduction

This paper presents an analytic method for computing sun and sky illumination under a canopy of trees. Recent illumination models have become increasingly accurate in accounting for light coming from the whole environment, and have also become increasingly expensive to compute.

Distributed ray tracing [1] requires enough rays at each pixel to adequately sample the illuminating environment. For a point light source, a single 'shadow feeler' will suffice, but for area light sources, multiple rays must sample the source area, and for diffuse reflection of the environment, the rays must sample the whole hemisphere above the surface position [2]. Certain simplifications result if the environment is 'at infinity', so that it is independent of the viewer's position [3,4,5]. A clear or cloudy sky is, for practical purposes, at infinity.

Other polygon-based algorithms generate the environment analytically, with a visible-surface computation from the point of view of each surface position. Finding the visible

surfaces of only the light sources is sufficient to render shadow penumbras from area light sources [8,19]. The ‘radiosity’ method can handle diffuse [6] and specular[7] interreflection between surfaces. In order to be computationally practical, these analytic algorithms compute the environment at only a sampled collection of surface points, and render the image by interpolation. Since the Fast Fourier Transform techniques used here compute all the shadows simultaneously, interpolation is not required.

Nishita and Nakamae [9] compute sky illumination using semicircular scan lines across a sky hemisphere with the CIE standard sky luminance functions [10]. But it is impractical to make these scan lines fine enough to sample the sun, so a separate shadow computation is required to include direct sunlight. The method presented here treats the sky as a raster image, so it can include skylight, cloud glow, and sunlight in a uniform manner. The shadow casting objects are also represented in a raster image, allowing finely detailed leaf shadows. However the method assumes all shadow casting objects lie near a canopy plane parallel to the ground, and only works for shadows cast onto the ground. Shadows cast onto other surfaces must be computed separately by some other method.

Shadows as Correlations

Consider the shadows cast onto the ground plane G in figure 1, by a collection of polygonally defined leaves and branches lying near a canopy plane C parallel to the ground, from a point light source at infinity in the direction L . The shadows can be considered as a polygonally defined surface texture on the plane G . If this texture is translated in the direction L onto the plane C , it defines a collection Q of polygons in the plane C which would generate the same shadows as the actual leaf and branch polygons. We will compute the sun, sky, and cloud illumination, as obscured by the polygons Q in C . Since rays from the sun are all nearly parallel to L , and the leaves all lie near C , the errors in penumbras cast by the leaves onto the ground will be of second order. (See caption to figure 1.) The illumination from directions farther away from L is much less bright, so the

errors in assuming all the leaves lie in C will not be very noticeable.

The polygons in the set Q on the canopy plane C can be scan converted into an anti-aliased raster transparency mask, $m(i, j)$, which is 1 for unobscured pixels, 0 for totally shadowed pixels, and some value between 0 and 1 for partially obscured pixels. Since it is this mask raster which will generate the shadows, polygons are not actually required. Any method which can generate a raster shadow texture on the ground plane (for example, a visible surface algorithm from the point of view of the sun) can be used to prepare the transparency mask.

Let $E(L)$ be the sky color in the direction L . Instead of projecting the sky environment onto a hemisphere at infinity, we project it onto a sky plane at infinity, parallel to the ground. The sky will appear the same from any point on the ground; its color will depend only on the direction of a ray, not the position. Thus we can think of the sky as an infinite plane S tangent to the unit hemisphere H of upward directions. As shown on figure 2, an element of solid angle $d\omega = d\theta(\sin\theta d\phi)$ on the hemisphere H corresponds to an area $dA = (\sec^2\theta d\theta)(\tan\theta d\phi)$ on the plane S , so $d\omega = \cos^3\theta dA$. Then, when there are no shadows, the unobscured intensity of the ground is (see [5])

$$\begin{aligned} I_0 &= \int_H f_r(V, L) E(L) \cos\theta d\omega \\ &= \int_S f_r(V, L) E(L) \cos^4\theta dA, \end{aligned} \tag{1}$$

where $f_r(V, L)$ is the bidirectional reflection function of the viewing direction V and the illumination direction L .

In the images presented here, I used Lambert's law, where $f_r(V, L)$ is constant. But the method works whenever $f_r(V, L)$ is independent of V , so it is appropriate whenever the pyramid of view containing the possible viewing vectors V is small, and f_r varies slowly. Because of the $\cos^4\theta$ weighting factor, it is possible to get a good approximation to I by integrating only over a square region T of the infinite plane S , as long as T contains the

sun. On T , we represent $f_r(V, L)E(L)\cos^4\theta$ as an illumination raster $e(i, j)$, defined for $0 \leq i, j \leq n-1$. Then $I_0 = \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} e(i, j)$, up to a normalizing factor which we will neglect here for simplicity.

Now we wish to include the shadows, and to compute an intensity raster $g(i, j)$ at each point on the ground plane G . We organize the coordinate systems on the mask raster $m(i, j)$, the illumination raster $e(i, j)$, and the ground intensity raster $g(i, j)$, as shown on figure 3. We assume the leaves are dark and opaque, and contribute no illumination of their own; they merely obscure light from the sky. A ray from $(0, 0)$ on the ground, in the direction (i, j) on the sky plane, will intersect the canopy plane at index (i, j) also, so $g(0, 0) = \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} e(i, j)m(i, j)$.

Now extend $m(i, j)$ and $g(i, j)$ to be periodic functions in i and j , of period n . This is equivalent to assuming that the canopy of trees is periodic, which is reasonable for a section of a tightly packed forest. For an isolated clump of trees, we scale the rasters so there are gaps between the periodic clumps, and then artificially remove the extraneous shadows, by assuming points on the ground outside the first period of $g(i, j)$ are completely unshadowed. Because of the $\cos^4\theta$ factor in equation (1), the errors introduced by the periodicity rapidly become small if the periods are made large enough.

As shown in figure 3, a ray from (k, l) on the ground, in the direction (i, j) , will intersect the canopy at $(i + k, j + l)$ so

$$g(k, l) = \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} e(i, j)m(i + k, j + l). \quad (2)$$

This means that g is the *correlation* of e and m . (See [11].)

The summation in equation (2) requires n^2 multiplications for each pixel (k, l) , and there are n^2 pixels, so the total computation is of order n^4 . With Fast Fourier Transform (FFT) techniques, it is possible to compute the correlation at all pixels in time of order $n^2 \log n$,

as follows. Let G , E , and M be the fourier transforms of g , e , and m respectively. Then by the correlation theorem (see [11]), $G = E\bar{M}$, where \bar{M} is the complex conjugate of M . Therefore, to get g , we use the FFT algorithm to find E and M , multiply E by the complex conjugate of M at each of the n^2 points in frequency space, and then use the FFT algorithm to find the inverse fourier transform g of $E\bar{M}$. On the Cray-1 or the Cray XMP, I used an FFT algorithm for two dimensional arrays, which transforms in one dimension while vectorizing over the other. In addition, since e and m are real, the FFT along the first index gives only $n/2 + 1$ independent complex numbers for each column of pixels, instead of n , resulting in a time saving factor of almost 2.

Implementation and Results

The shadows of the leaves and branches on each other were computed by the method of Max [12] along radial scan lines. Each scan line is the intersection with the picture plane of a scan plane through the eye and light source. After all the shadow segments were accumulated in each scan plane, they were written out into a file. The scan planes also intersect the canopy plane C in a series of radial scan lines. Therefore, the shadows in the file can be scan converted into the mask raster $m(i, j)$ shown in figure 4, using the anti-aliasing algorithm in Max [13].

The image of the cloudy sky shown in figure 5 was generated by the method of Gardner [14]. The sun was added, with the sun glow through the clouds calculated by the method of Blinn [15]. When multiplied by $\cos^4 \theta$, this resulted in the raster $e(i, j)$ shown in figure 6. A raster $g(i, j)$ for the intensity reflected from the environmental illumination by the ground, as shown in figure 7, was computed by the FFT correlation technique discussed above. Then the ground and sky images were resampled in perspective as a background raster in the picture plane, using standard area averaging techniques [16,17].

The visible segments of the trees were also scan converted as a foreground raster, shown in

figure 8, using the methods of [13]. An α channel, shown in figure 9, carrying pixel coverage information [18], was produced at the same time. Then the background and foreground were combined into the final image, shown in figure 10:

$$final(i, j) = foreground(i, j) + (1 - \alpha(i, j))background(i, j).$$

If atmospheric illumination is desired (see [12]), it can be scan converted as an extra channel, and added on top of the final image.

Note the effects of the sky illumination in figures 7 and 10. The ground is darker in the center of the shadow of a large clump of trees than it is in shadows of branches or smaller clumps. Also, the ground is brighter far from the shadow of the trees than it is nearer to the shadow, but still outside its penumbra from the sun. Compare this with figure 11, without these effects and without penumbras. Figures 12 and 13 show the effects of a partial eclipse on the penumbra of the sun.

To produce computer animation showing the illumination change in a grove of trees as the moon or clouds cover the sun, I used a short cycle of images of the trees, without ground or sky, showing the leaves fluttering in the breeze. I rendered the foreground images twice, once including point source illumination at the sun, and once with only diffuse illumination from the sky. I also computed the shadow mask raster $m(i, j)$ for each frame of the cycle, and its fourier transform. Then for each frame of the longer final sequence, the background image is found as above, using the precomputed fourier transform M of the mask m .

The fraction S of the direct sunlight which gets through the clouds is used to interpolate between the two foreground images, which are then combined with the background using the precomputed α channel. Finally, the atmospheric illumination is multiplied by S , and added to the image.

Conclusions and Extensions

Image correlation by Fast Fourier Transforms is an efficient technique for computing shadows on the ground from environmental illumination, provided the shadow casting objects lie in or near a single horizontal plane. Mask rasters m_1 and m_2 , in two planes at different heights, combine with the sky raster e in a more complex formula, for which there is no obvious *FFT* speedup. However, if we assume the two masks are uncorrelated, we can approximate the composition of their two shadowing effects by computing the two ground intensity rasters g as above, multiplying them together, and dividing by the unobstructed intensity I_0 .

Acknowledgements

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract number W-7405-ENG-48. I wish to thank Brian Cabral, Becky Springmeyer, Chuck Grant, Michael Gwilliam, and Jeff Kallman for helpful suggestions on the exposition, Becky Springmeyer for the typing and layout, and Brian Cabral for providing the cloud texture in figure 5.

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, processes, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or other product endorsement purposes.

Figures

Figure 1. The canopy plane C parallel to the ground plane G . The light source is in the direction L . Polygons P_1 , and P_2 near C cast shadows S_1 and S_2 on G , which translate to Q_1 and Q_2 in C . The positioning error e in the shadow of P_1 from a direction L' , caused by using Q_1 instead, varies as the product of the distance d of P_1 from C and the angle θ between L' and L .

Figure 2. The unit hemisphere H of directions above a surface point O , and the plane S tangent to it at the north pole N . An element of solid angle on H , with sides $d\theta$ and $\sin\theta d\phi$, projects to an element of area on S with sides $\sec^2\theta d\theta$ and $\tan\theta d\phi$.

Figure 3. The coordinates on the canopy plane C are scaled to correspond to those on the plane S of directions above the origin $(0,0)$ on the ground G . The coordinates on the ground plane G are vertically translated from those on C . When the hemisphere H and plane S of directions are moved above the pixel (k,l) on G , the ray in direction (i,j) intersects C in the point $(i+k, j+l)$.

Figure 4. Transparency mask $m(i,j)$ from leaf shadows.

Figure 5. Image of cloudy sky.

Figure 6. Sky raster $e(i,j)$. The small sun disc in the center is actually one thousand times as bright as the brightest clouds.

Figure 7. Ground intensity raster $g(i,j)$.

Figure 8. Trees, with shadows computed by the method of Max [12], as if the sun were a point source.

Figure 9. The α channel, used for compositing

Figure 10. The final image, made by resampling figures 5 and 7 in perspective, and compositing figure 8 on top.

Figure 11. Similar to figure 10 but without sky illumination and penumbras, using figure 4 instead of figure 7.

Figure 12. The ground raster for an eclipsed sun. Note crescent shaped spots of light.

Figure 13. A composite image in a partial eclipse, using figure 12 as the ground raster.

References

- 1 Cook, Robert L., Porter, Thomas, and Carpenter, Loren. *Distributed Ray Tracing*. Computer Graphics Vol. 18, No. 3. (1984), pp. 137-145.
- 2 Kajiya, James T. *The Rendering Equation*. Computer Graphics Vol. 20, No. 4. (1986), pp. 143-150.
- 3 Greene, Ned. *Applications of World Projections*. IEEE CG&A Vol. 6, No. 11. (1986), pp. 21-29.
- 4 Miller, Gene S. and Hoffman, Robert C. *Illumination and Reflection Maps: Simulated Objects in Simulated and Real Environments*. SIGGRAPH '84 Advanced Computer Graphics Animation Course Notes, July (1984).
- 5 Cabral, Brian, Max, Nelson, and Springmeyer, Rebecca. *Bidirectional Reflection Functions from Surface Bump Maps*. Submitted to SIGGRAPH '86.
- 6 Cohen, Michael F., and Greenberg, Donald D. *A Radiosity Solution for Complex Environments*. Computer Graphics. Vol. 19, No. 3. (1985), pp. 31-40.
- 7 Immel, David S., Cohen, Michael F., and Greenberg, Donald P. *A Radiosity Method for Non-Diffuse Environments*. Computer Graphics Vol. 20, No. 4. (1986), pp. 133-142.
- 8 Nishita, Tomoyuki and Nakamae, Eihachiro. *Half Tone Representation of 3-D Objects Illuminated by Area Sources or Polyhedron Sources*. IEEE Proc. of COMSAC (1983), pp. 237-241.
- 9 Nishita, Tomoyuki and Nakamae, Eihachiro. *Continuous Tone Representation of Three*

- Dimensional Objects Illuminated by Sky Light*. Computer Graphics Vol. 20, No. 4. (1986), pp. 125-132.
- 10 CIE Technical Committee 4.2: Standardisation of Luminance Distribution on Clear Skies. *CIE Publication No. 22*. Commission International de l'Eclairage. Paris, (1973), p. 7.
 - 11 Gonzalez, Rafael C. and Wintz, Paul. *Digital Image Processing*. Addison Wesley. Reading, Mass. (1977).
 - 12 Max, Nelson. *Atmospheric Illumination and Shadows* Computer Graphics. Vol. 20, No. 4. (1986), pp. 117-124.
 - 13 Max, Nelson. *Anti-Aliasing Scan Line Data*. submitted to IEEE Computer Graphics and Applications.
 - 14 Gardner, Geoffrey Y. *Visual Simulation of Clouds*. Computer Graphics Vol. 19, No. 3. (1985), pp. 297-303.
 - 15 Blinn, James F. *Light Reflection Functions for Simulation of Clouds and Dusty Surfaces*. Computer Graphics Vol. 16, No. 3. (1982), pp. 21-29.
 - 16 Catmull, Ed and Smith, Alvy Ray. *3-D Transformations in Scanline Order*. Computer Graphics Vol. 14, No. 3. (1980), pp. 279-285.
 - 17 Fant, Karl M. *A Nonaliasing, Real Time Spatial Transform Technique*. IEEE Computer Graphics and Applications Vol. 6, No. 1. (1986), pp. 71-80. (See also Letters to the Editor in Vol. 6, Nos. 3 and 7).
 - 18 Porter, Thomas and Duff, Tom. *Compositing Digital Images*. Computer Graphics Vol. 18, No. 3. (1984), pp. 253-259.
 - 19 Joy, Kenneth. *Soft Shadow Volumes*. Submitted to SIGGRAPH '87.

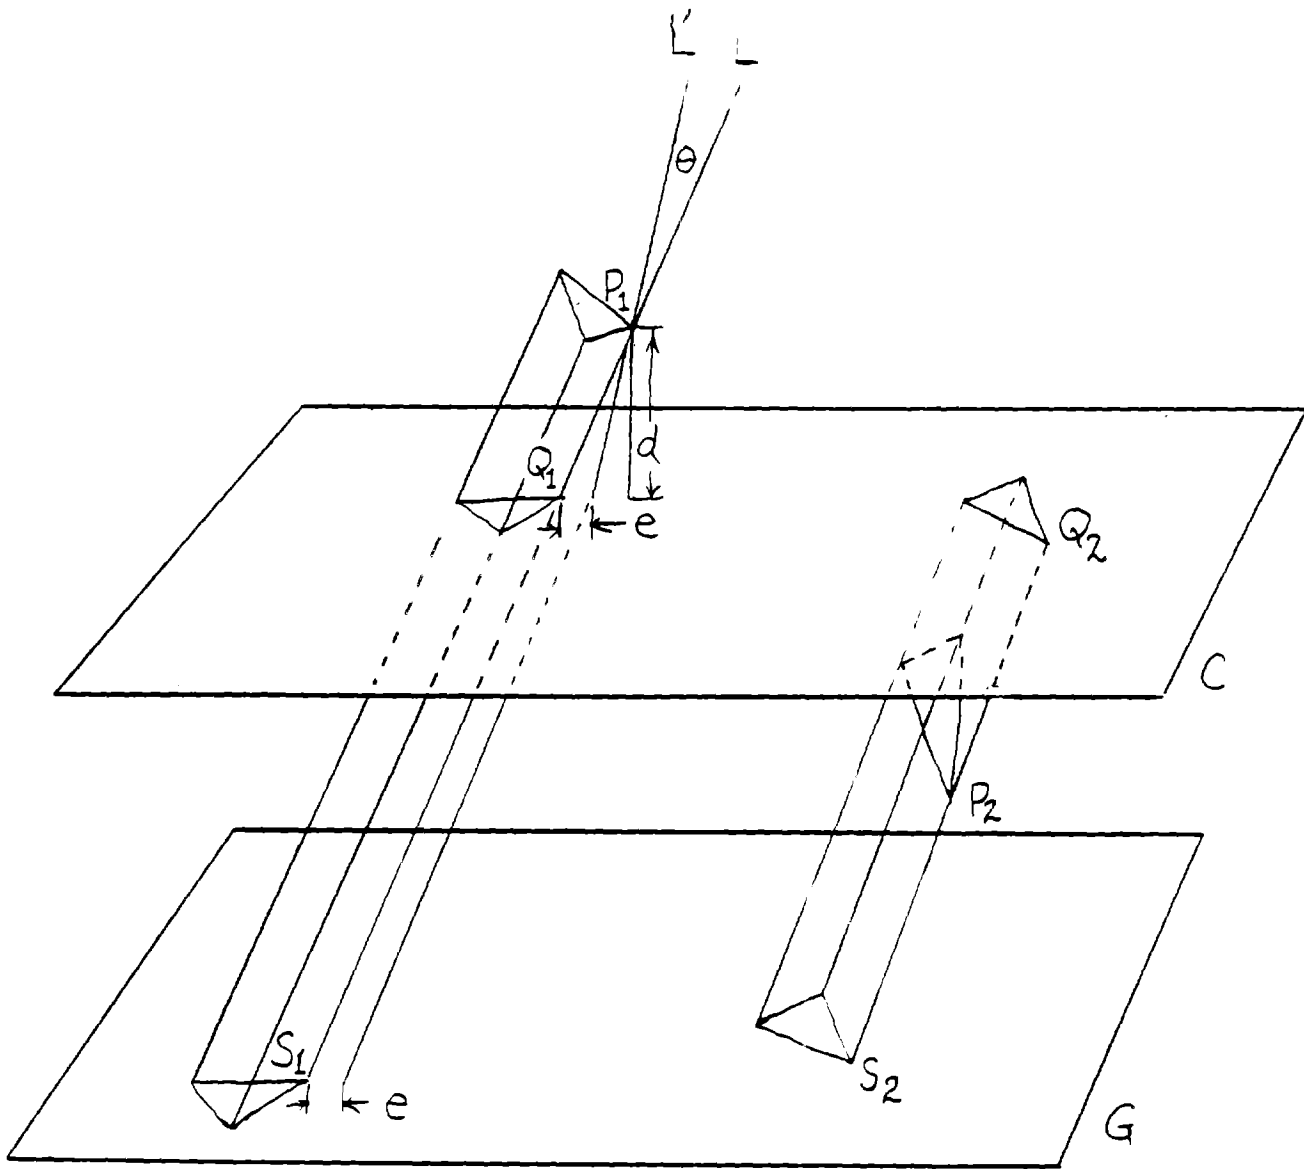


Figure 1. The canopy plane C parallel to the ground plane G . The light source is in the direction L . Polygons P_1 , and P_2 near C cast shadows S_1 and S_2 on G , which translate to Q_1 and Q_2 in C . The positioning error e in the shadow of P_1 from a direction L' , caused by using Q_1 instead, varies as the product of the distance d of P_1 from C and the angle θ between L' and L .

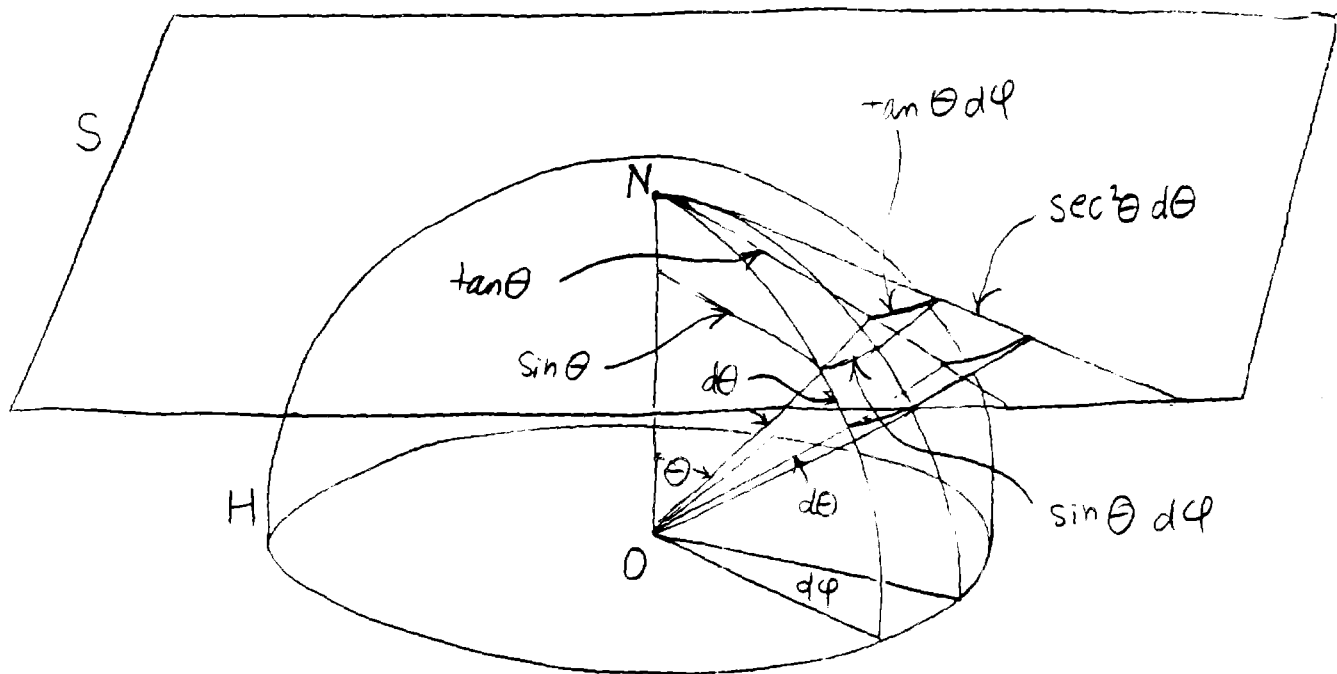


Figure 2. The unit hemisphere H of directions above a surface point O , and the plane S tangent to it at the north pole N . An element of solid angle on H , with sides $d\theta$ and $\sin \theta d\phi$, projects to an element of area on S with sides $\sec^2 \theta d\theta$ and $\tan \theta d\phi$.

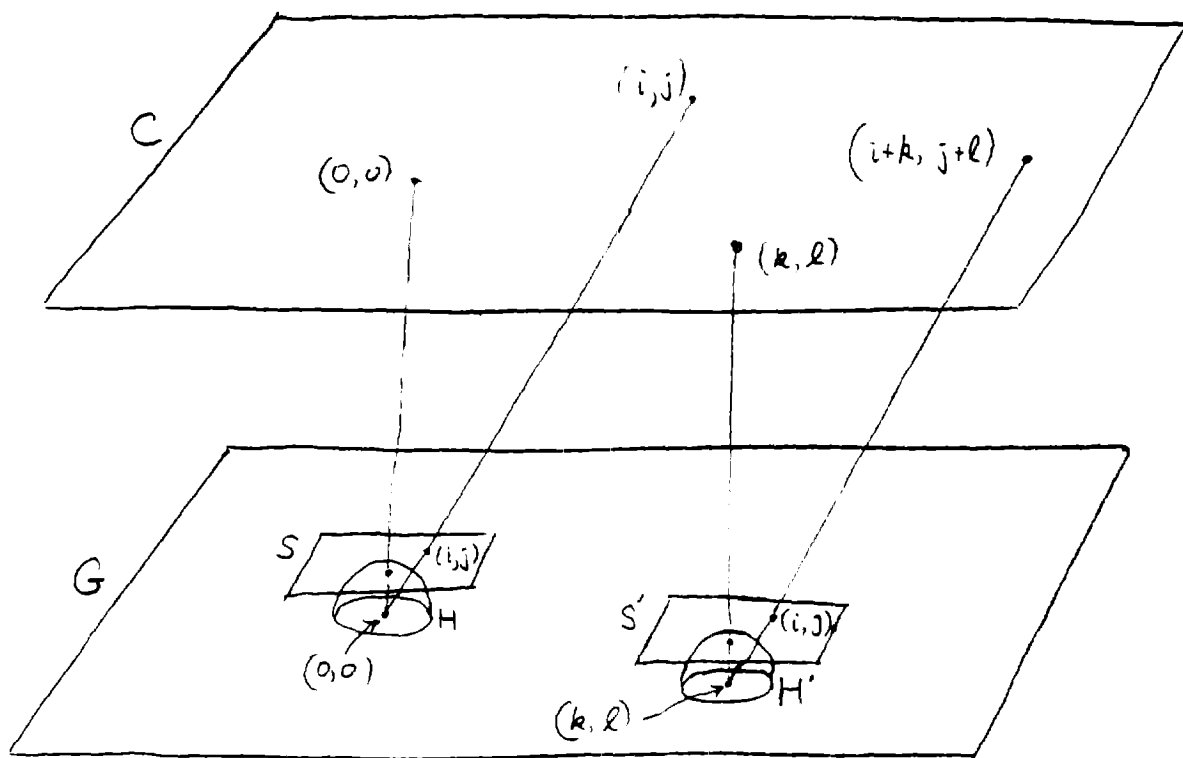


Figure 3. The coordinates on the canopy plane C are scaled to correspond to those on the plane S of directions above the origin $(0,0)$ on the ground G . The coordinates on the ground plane G are vertically translated from those on C . When the hemisphere H and plane S of directions are moved above the pixel (k,l) on G , the ray in direction (i,j) intersects C in the point $(i+k, j+l)$.



